Article

The Representation Economy: A Framework for AI, Institutions, and Machine-Legible Reality

Raktim Singh

Independent Researcher, Bangalore, India; e-mail: raktims2210@gmail.com; ORCID: 0009-0002-6207-602X

Correspondence: raktims2210@gmail.com

Author note: This paper is written in a personal capacity. It introduces an original conceptual framework for understanding AI-mediated value creation, institutional governance, and enterprise execution. It does not represent the views of any employer or institution.

Abstract

Artificial intelligence is increasingly embedded in enterprises, markets, public services, and institutional decision-making. Yet most AI discourse still centers on models, algorithms, automation, and compute while underestimating the representation systems that make reality legible to machines. This paper introduces the Representation Economy as an emerging economic and institutional condition in which value, risk, power, and trust are shaped by the quality and control of machine-legible representations of entities, events, relationships, states, and rights. Using conceptual systems analysis and enterprise architecture reasoning, the paper develops the SENSE-CORE-DRIVER framework as an institutional architecture for AI. SENSE is the legibility layer through which signals, entities, states, and change become machine-readable. CORE is the cognition layer through which models, reasoning systems, and agents interpret context and generate decisions. DRIVER is the legitimacy and execution layer through which delegation, representation, identity, verification, execution, and recourse convert decisions into accountable action. The paper argues that many enterprise AI failures are not only model failures but also failures of representation and legitimacy. It concludes with implications for AI governance, autonomy allocation, enterprise architecture, representation quality engineering, DRIVEROps, and future research.

Keywords: Representation Economy; SENSE-CORE-DRIVER; Enterprise AI; AI Governance; Machine-Legible Reality; Institutional AI; Intelligent Systems; AI Architecture; Trustworthy AI; Institutional Intelligence

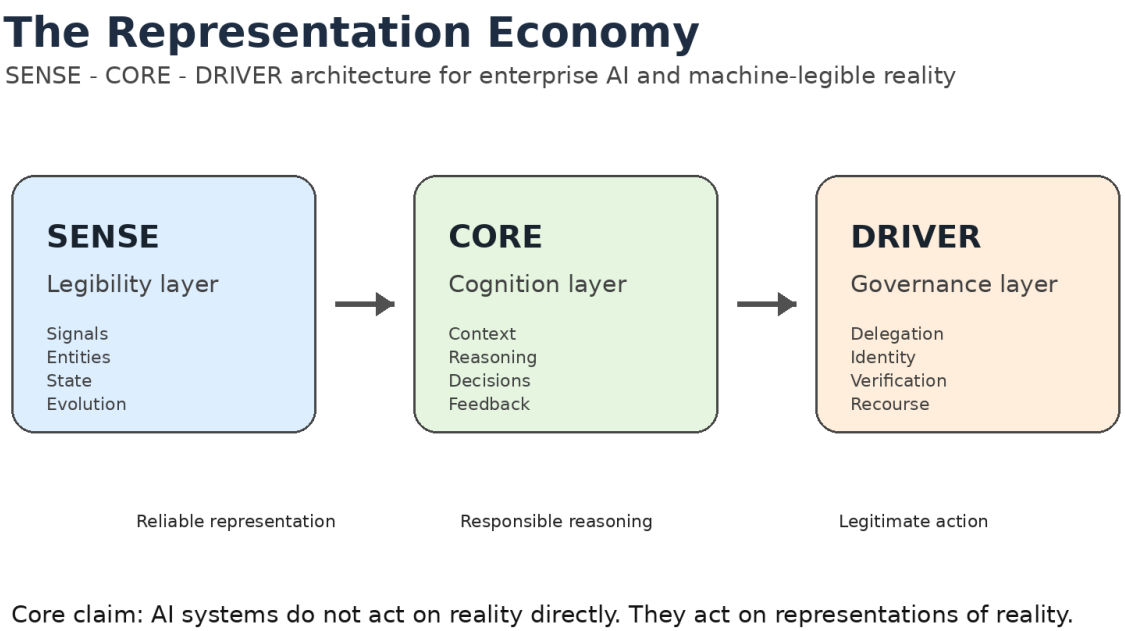

Figure 1. Graphical abstract: Representation Economy and SENSE-CORE-DRIVER.

1. Introduction: AI Does Not Act on Reality Directly

The dominant narrative of artificial intelligence has been organized around intelligence itself. Public discussion asks whether models are more capable, whether agents can reason, whether automation will replace work, and whether compute scale will continue to generate better performance. These questions matter, but they leave out a deeper institutional problem: AI systems do not act on reality directly. They act on representations of reality.

A bank does not lend to a person in the abstract. It acts on represented income, repayment history, identity attributes, fraud signals, collateral data, policy constraints, and authorization flows. A hospital does not treat a person only as a living body. It acts on records, test results, symptoms, images, triage categories, clinical guidelines, and coverage rules. A supply chain system does not see the world directly. It sees orders, sensors, invoices, inventory positions, shipment states, exceptions, and commitments. In each case, the machine-readable representation becomes the operational reality for the system.

This is why the next stage of the AI economy cannot be understood only as a competition over models. It must also be understood as a competition over representation: who can sense the world with fidelity, attach signals to the right entities, maintain state over time, reason on top of that state, and act with legitimate authority. The firms and institutions that solve this problem will not merely use AI. They will shape the categories, boundaries, and permissions through which AI acts.

The central claim of this paper is that the next AI advantage will not come only from better intelligence. It will come from better representation and more legitimate execution. This claim extends earlier work on prediction economics [1], the sciences of the artificial [2], information ethics [3], institutional economics [4,5], and digital governance [6]. It also responds to the practical pressure now facing enterprises as AI moves from answering questions to influencing workflows, routing exceptions, triggering actions, supporting employees, generating code, managing operations, and recommending interventions.

Once AI begins to influence institutional decisions, the central question changes. It is no longer only: can the AI produce an intelligent answer? It becomes: can the institution sense reality correctly, reason over it responsibly, and act with legitimate authority? The SENSE-CORE-DRIVER framework addresses that question. It provides

a practical language for CIOs, CTOs, architects, risk leaders, boards, regulators, and product teams to discuss the full system around AI, not only the model inside it.

2. Methodological Approach

This paper is a conceptual and architectural article rather than an empirical experiment. It develops a framework through systems analysis, enterprise architecture interpretation, institutional governance reasoning, and synthesis across AI governance, decision systems, digital representation, platform economics, and sociotechnical systems. The purpose is not to claim a predictive scientific law. It is to offer a practical conceptual architecture that can help organizations diagnose, design, and govern AI-mediated systems.

The approach has four steps. First, it distinguishes data from representation. Second, it identifies representation as an economic and institutional layer. Third, it separates AI-enabled institutional systems into three functional layers: SENSE, CORE, and DRIVER. Fourth, it derives enterprise, governance, market, and research implications from the interaction among these layers.

The framework is informed by the observation that AI systems increasingly operate inside institutional processes rather than isolated technical environments. In production, AI is rarely just a model. It is a system of signals, entities, state, context, models, policies, identities, permissions, actions, controls, and recourse paths. A framework that analyzes only the model cannot explain the full risk or value of such systems.

3. The Representation Economy: Definition and Scope

The Representation Economy is the emerging condition in which economic value, institutional power, and social trust are increasingly shaped by how entities, events, states, relationships, and rights are represented for machine interpretation and action. In earlier digital eras, organizations competed over data collection, software distribution, workflow control, platform access, and network effects. In the AI era, these sources of advantage do not disappear. They are reorganized around machine-legible representation.

AI systems require representations that are current, contextual, structured, authorized, and actionable. The scarce asset is not simply data volume. It is the ability to convert messy reality into usable institutional state. A transaction record is data. A continuously updated risk profile is representation. A sensor reading is data. A machine-state model that determines when maintenance should occur is representation. A customer chat transcript is data. A structured account of intent, eligibility, unresolved issues, prior commitments, and permissible next actions is representation.

The distinction matters because AI value depends less on isolated data points and more on the structure that gives those points meaning. A model can generate a response from weak context, but an institution cannot responsibly delegate action to a system that lacks a reliable representation of the entity it is affecting.

Representation is becoming an economic layer because it determines what AI systems can perceive, decide, and execute. In platform markets, power emerged from controlling interfaces, ecosystems, and participation rules [7]. In the Representation Economy, power emerges from controlling the machine-legible models of reality on which decisions depend. This creates new strategic questions. Who owns the representation of a customer, asset, product, worker, risk, city, supply chain, or ecosystem? Who updates it? Who verifies it? Who has the right to contest it? Who benefits when the representation becomes the default interface through which AI systems act? These are technical questions, but they are also questions of governance, markets, and institutional legitimacy.

Representation is not neutral. The act of encoding reality into machine-readable form creates categories, priorities, omissions, confidence levels, and authority boundaries. A person becomes a profile. A factory becomes sensor streams. A supplier becomes a risk score. A city becomes traffic feeds. A company becomes a knowledge graph. A legal obligation becomes policy logic. A conversation becomes embeddings. A complaint becomes a category. A transaction becomes a signal. The design of that translation influences what AI can see and what institutions can justify.

4. The SENSE-CORE-DRIVER Framework

To analyze this new layer, this paper proposes the SENSE-CORE-DRIVER framework. The framework separates an AI-enabled institutional system into three layers: the legibility layer, the cognition layer, and the legitimacy-and-execution layer. SENSE is how reality becomes machine-legible. CORE is how intelligence interprets, reasons, decides, and learns. DRIVER is how decisions become authorized, governed, verified, executed, and corrected.

In traditional enterprise technology, these layers were often hidden inside applications, workflows, reports, business rules, and human judgment. In the AI era, they must become explicit. Enterprise AI does not fail only because the model is weak. It fails when the institution has not clearly designed what the system can see, what it is allowed to infer, who authorized it to act, how its action will be verified, and what happens when it is wrong.

SENSE **Legibility layer**	**CORE** **Cognition layer**	**DRIVER** **Legitimacy and execution layer**
Signal detection	Comprehend context	Delegation
Entity linkage	Optimize decisions	Representation rights
State representation Evolution over time	Realize decisions Evolve through feedback	Identity Verification Execution Recourse

Figure 2. SENSE-CORE-DRIVER separates AI systems into legibility, cognition, and legitimacy/execution layers.

4.1 SENSE: The Legibility Layer

SENSE is the layer where reality becomes machine-legible. It answers the question: what does the system believe reality is? SENSE includes Signal, ENtity, State representation, and Evolution. Signal means detecting relevant events, changes, anomalies, traces, documents, actions, and environmental inputs. ENtity means attaching signals to a persistent actor, object, location, asset, account, process, or institutional unit. State representation means building a structured model of the current condition, rights, obligations, risks, constraints, and context of that entity. Evolution means updating that state over time as new signals arrive, old assumptions decay, and relationships change.

SENSE is often underestimated because it looks like data engineering. But it is not only data engineering. It is institutional representation design. It determines which events matter, which entities exist, which relationships are visible, which states are trusted, and which changes are propagated downstream. If SENSE is weak, CORE may reason over conflicting or stale reality, and DRIVER may execute a wrong action with formal confidence.

Consider a delivery address problem. A customer changes an address in one system but not another. The warehouse has the old address. The payment system has a billing address. The delivery partner has a partial address. Customer support has a chat message saying to deliver to an office this week. If an AI agent is asked where the package should be delivered, this is not mainly a model reasoning problem. It is a SENSE problem. The system must know which address is current, which applies to this order, which instruction is authorized, which source is authoritative, whether delivery has already been dispatched, and whether the address can still be changed.

In the AI economy, enterprises will not compete only on access to better intelligence. They will increasingly compete on who can create the most trusted, current, interoperable, and actionable representation of reality. Traditional data quality asks whether data is accurate, complete, consistent, and timely. SENSE asks a deeper question: is the institution representing reality well enough for intelligent action?

4.2 CORE: The Cognition Layer

CORE is the cognition layer. It includes model reasoning, planning, optimization, prediction, generation, classification, summarization, simulation, decision support, and agentic planning. CORE is where AI capability is most visible and where most public excitement concentrates. It answers the question: how does the system interpret what it senses?

CORE can be described through four functions: Comprehend context, Optimize decisions, Realize action plans, and Evolve through feedback. Large language models, multimodal models, prediction systems, optimization engines, retrieval-augmented systems, and agentic planning systems largely sit in this layer. CORE is powerful, but it is not sovereign. It depends on SENSE for grounding and on DRIVER for legitimacy.

Many organizations fall into the intelligence trap. The intelligence trap occurs when organizations assume that better reasoning will compensate for poor institutional architecture. It will not. A better model cannot fix a broken representation of reality. A more autonomous agent cannot create legitimate authority by itself. A more accurate prediction cannot decide whether an action should be allowed. A faster recommendation cannot explain who is accountable.

For architects, CORE should not be imagined as one model. Modern enterprise CORE may include a retrieval layer that brings enterprise context, a model layer that interprets the question, a planning layer that breaks down work, a tool layer that interacts with systems, a policy layer that constrains possible actions, a simulation layer that tests likely outcomes, a verification layer that checks results, and a feedback layer that improves future behavior. As CORE becomes more capable, SENSE and DRIVER become more important. Weak AI can be contained by limited capability. Strong AI requires stronger institutional design.

4.3 DRIVER: The Legitimacy and Execution Layer

DRIVER is the layer where intelligence becomes legitimate action. It answers the question: who allowed the system to act, on what basis, under what constraints, with what evidence, and with what path for correction? DRIVER includes Delegation, Representation, Identity, Verification, Execution, and Recourse.

Delegation defines who authorized the system to act, under what scope, and with what boundary conditions. Representation defines what model of reality the system used and whether that model is appropriate for the decision. Identity defines which person, asset, account, process, or institution is affected. Verification defines how the decision is checked before, during, or after execution. Execution defines how the action is carried out through enterprise systems, workflows, APIs, humans, or agents. Recourse defines what happens if the action is wrong, harmful, disputed, or no longer valid.

DRIVER matters because intelligence does not create legitimacy. A system may produce a correct recommendation and still lack authority to act. Conversely, a system may be authorized to act but lack sufficient verification and recourse. As AI systems become more autonomous, the DRIVER layer becomes a central determinant of trust.

Imagine a retail company using AI to approve refunds. A small refund for a verified damaged item may be safe to automate. A large refund for a high-value product may need human approval. A repeated refund pattern may require fraud review. A legally sensitive case may require compliance review. DRIVER defines the action boundary. It tells the system when it may recommend, when it may execute, when it must escalate, when it must not act, what evidence it must preserve, and how recourse must be provided. Without DRIVER, AI governance becomes fragile. The system may be accurate most of the time but illegitimate when it matters.

5. Why AI Failures Are Often Representation Failures

Much of the discussion around AI failure focuses on hallucination, bias, model weakness, prompt design, or insufficient training data. These are real issues, but they are often symptoms of a deeper pattern. The system failed because it represented the world incorrectly, incompletely, or illegitimately.

In enterprise settings, AI outputs are embedded in business processes. A system may misroute an email not because the model cannot read language, but because it lacks a faithful representation of process ownership, customer context, exception rules, or contractual obligations. A code assistant may generate plausible code but fail to represent architectural constraints, security policies, dependency histories, or deployment environments. A customer-support agent may summarize a problem correctly but act incorrectly because it lacks authority boundaries and recourse mechanisms.

The SENSE-CORE-DRIVER framework reframes AI failure as a propagation problem. Weak SENSE feeds uncertain CORE. Uncertain CORE produces fragile decisions. Fragile decisions enter DRIVER without sufficient authority, verification, or reversibility. The visible failure may appear as a model failure, but the origin may be representational or institutional.

The major failure modes are layer-specific. A SENSE failure occurs when the system detects the wrong signal, misses the relevant entity, maintains stale state, or ignores change over time. A CORE failure occurs when the system reasons, predicts, or generates on top of incomplete or misleading representation. A DRIVER failure occurs when the system acts without legitimate delegation, clear accountability, reliable verification, or meaningful recourse. A propagation failure occurs when weaknesses in one layer are amplified by the next layer, producing institutional risk that appears as an AI problem but originates as a representation problem.

6. SENSE-CORE-DRIVER as Enterprise AI Control Architecture

The practical value of SENSE-CORE-DRIVER is that it gives enterprises a way to structure AI systems before they scale. It helps answer three architectural questions: what can the AI see, how does the AI reason, and what is the AI allowed to do? These questions sound simple, but they are the foundation of enterprise AI governance.

A customer service agent can be mapped through the framework. SENSE includes customer identity, history, product context, open tickets, sentiment, and contract terms. CORE includes intent recognition, response generation, next-best-action reasoning, and escalation reasoning. DRIVER includes permitted actions, refund limits, escalation rules, audit logs, and correction paths. A software engineering agent can also be mapped: SENSE includes codebase, dependencies, architecture, security rules, and issue history. CORE includes code generation, test creation, refactoring logic, and debugging reasoning. DRIVER includes repository permissions, approval workflow, security verification, and rollback.

This is why SENSE-CORE-DRIVER is useful across industries. It does not replace existing governance frameworks. It helps operationalize them. NIST AI RMF emphasizes governance, mapping, measurement, and management of AI risks [11]. ISO/IEC 42001 provides a management-system approach for organizational AI governance [12]. The EU AI Act follows a risk-based approach to AI regulation [13]. OECD AI principles emphasize trustworthy, human-centered AI [14]. SENSE-CORE-DRIVER provides a design lens that can help connect these governance goals to system architecture.

Use case	SENSE	CORE	DRIVER
Customer service agent	Customer identity, history, contract, open tickets	Intent recognition, next best action, response generation	Refund limits, escalation, audit log, correction path
Software engineering agent	Codebase, dependencies, architecture, security rules	Code generation, test creation, debugging, refactoring	Repository permissions, approvals, security checks, rollback
Finance risk system	Transactions, exposure, market signals, counterparties	Anomaly detection, risk scoring, scenario analysis	Approval authority, compliance constraints, evidence retention
Manufacturing AI system	Machine telemetry, maintenance history, quality signals	Failure prediction, root cause analysis, optimization	Shutdown authority, safety rules, recovery procedure

Table 1. Example mapping of enterprise AI use cases through SENSE, CORE, and DRIVER.

7. Autonomy Allocation: When Should AI Act?

One of the most important enterprise AI decisions is autonomy allocation. Not every process needs an AI agent. Not every decision should be automated. Not every workflow should use generative AI. Not every task needs reasoning. Not every exception should be delegated. SENSE-CORE-DRIVER helps determine the right level of autonomy.

If SENSE is stable, CORE ambiguity is low, and DRIVER risk is low, deterministic automation may be enough. An example is sending a standard notification after a verified event. If SENSE is strong, CORE ambiguity is moderate, and DRIVER has clear boundaries, AI-assisted recommendation may be appropriate. An example is suggesting next-best actions for a service agent. If SENSE is dynamic, CORE ambiguity is high, and DRIVER consequences are serious, human judgment must remain central. An example is resolving a complex customer dispute or approving a high-risk financial decision. If SENSE is weak, no amount of CORE intelligence should be trusted with action.

The framework therefore does not say use AI everywhere. It says allocate autonomy based on representation quality, reasoning need, and legitimacy risk. This moves enterprises away from random AI adoption and toward governed AI deployment.

8. Representation Quality Engineering

Enterprise AI requires a new discipline: representation quality engineering. Software quality engineering traditionally focused on whether applications behave as expected. Model evaluation focuses on model performance. Representation quality engineering asks whether the system is seeing the right reality before it reasons or acts.

Representation quality engineering asks whether signals are meaningful, entities are correctly linked, state is current, context is complete, relationships are accurate, exceptions are visible, updates are traceable, humans can correct the representation, and downstream AI systems can trust it. This discipline will become as important as model evaluation because a brilliant model can still fail when the reality provided to it is distorted.

For architects, representation quality engineering may involve data contracts, entity resolution tests, knowledge graph validation, provenance checks, semantic consistency checks, ontology governance, event freshness monitoring, state drift detection, and human correction workflows. For executives, the message is simpler: before you trust AI to decide, make sure it can see.

9. DRIVEROps: Operational Governance for AI Action

If SENSE needs representation quality engineering, DRIVER needs operational governance. This paper uses the term DRIVEROps to describe the operating discipline for delegation, verification, execution, recourse, and authority management in AI systems. DRIVEROps is where AI governance becomes operational rather than merely declarative.

DRIVEROps would manage who can delegate authority to AI systems, what actions AI can perform, which actions require approval, what evidence must be captured, which actions are reversible, which actions require explanation, how appeals are handled, how incidents are investigated, and how autonomy boundaries are changed over time. Many enterprises already have model risk management, cybersecurity governance, IT controls, privacy processes, and compliance reviews. AI agents create a new challenge: systems that can act across tools, workflows, and organizational boundaries. That requires runtime governance, not only design-time approval.

Runtime governance means the enterprise can control intelligent action while it happens. It is not merely a model card, policy document, committee review, or static checklist. It is an executable control layer over delegation, action, evidence, verification, and recourse.

10. Governance and Institutional Implications

AI governance often focuses on model risk, fairness, explainability, privacy, security, and compliance. These remain important. However, the Representation Economy adds another layer: representation governance. Institutions must govern not only what models do, but also what they are allowed to treat as reality.

Representation governance asks whether the system has correctly represented the relevant entity, whether the representation is current, whether affected parties can challenge it, whether the representation is traceable to evidence, and whether action taken on its basis is legitimate. Representation integrity is the degree to which a machine-legible representation is accurate, contextual, authorized, traceable, current, and contestable. It is not identical to data quality. Data quality can be high while representation integrity is low if the data is attached to the wrong entity, interpreted outside context, used beyond its authorized scope, or converted into action without recourse.

As representation becomes valuable, attacks will increasingly target the representation layer. These may include poisoning signals, manipulating entity linkage, creating false state histories, exploiting stale context, or gaming AI-visible profiles. Future security architectures will need to protect not only applications and data stores, but also the representations through which AI systems perceive reality.

The framework also clarifies why human oversight is necessary but insufficient. A human in the loop does not automatically create governance if the human cannot inspect the representation, understand the authority boundary, verify the evidence, reverse the action, or meaningfully intervene. Real governance requires designed intervention, not symbolic approval.

11. Market and Institutional Implications

The Representation Economy may create new categories of firms and institutional roles. Some companies will specialize in representation infrastructure: identity graphs, asset state systems, trusted data spaces, industrial twins, model context layers, verification services, provenance systems, and context assurance. Others may become representation utilities that provide trusted representations across ecosystems. Still others may build representation marketplaces where different parties contribute, validate, license, or contest machine-legible state.

At the same time, representation power can become concentrated. A company that controls the default representation of a customer, merchant, asset, city, or professional identity may influence how AI systems evaluate, recommend, price, approve, exclude, or act. This creates both strategic opportunity and governance risk.

Boards, regulators, and executives will need to ask who controls the representation layer, how representations are audited, whether affected entities have visibility and recourse, and whether AI-enabled decisions depend on contested or opaque representations. The legitimacy of AI systems will increasingly depend on the legitimacy of the representations they use.

12. Why This Is Not Another AI Ethics Framework

SENSE-CORE-DRIVER is not a replacement for AI ethics, law, standards, risk management, model evaluation, or software architecture. It is an institutional architecture that connects reality, intelligence, and authority. Ethics asks what should be right. Law defines what must be allowed. Risk management identifies what may go wrong. Technical architecture defines how systems work. SENSE-CORE-DRIVER connects these into an operating model for intelligent institutions.

Enterprise AI governance cannot be solved by principles alone. It must be translated into systems, roles, controls, evidence, workflows, and accountability. The framework provides that bridge. It is simple enough to remember, deep enough to design with, and practical enough to govern with.

The framework is also a translation layer between executive intent and technical implementation. A board may ask for transparency, but engineers need to know whether transparency means visibility into input signals, retrieved context, model reasoning, tool calls, authorization boundaries, or executed actions. A regulator may ask for accountability, but architects need to know whether accountability is attached to a model owner, workflow owner, data owner, process owner, human approver, or delegated agent. SENSE-CORE-DRIVER turns abstract governance language into inspectable system layers.

This matters because AI adoption is moving faster than institutional redesign. Many organizations deploy copilots, agents, and automation flows before they know how reality is represented, how authority is delegated, or how recourse will be provided. In that environment, governance can become symbolic. The framework pushes governance into architecture. It asks whether accountability is observable in the system itself, not merely declared in a policy.

13. A Diagnostic for Enterprise AI Readiness

Before scaling any AI use case, leaders can use nine questions. The SENSE questions are: what reality does this AI system depend on, how is that reality represented, and how do we know the representation is current and correct? The CORE questions are: what reasoning is the AI performing, what context does it use, and how do we verify the quality of its recommendation or action plan? The DRIVER questions are: who authorized the system to act, what is the boundary of action, and what happens if the system is wrong?

These questions force the enterprise to examine not just the model, but the full institutional system around the model. They also help leaders decide whether a use case should be automated, AI-assisted, human-led, or deferred until representation quality improves. The diagnostic is intentionally simple because enterprise AI governance must be usable by executives, architects, risk leaders, product owners, and operational teams.

The most important readiness question is not where can we add AI. The more mature question is where is SENSE strong enough, CORE useful enough, and DRIVER legitimate enough for AI to act. That question converts AI strategy from tool adoption into institutional design.

14. Research Agenda

The framework proposed in this paper is intended as a starting point. Several research directions follow. First, organizations need representation metrics: ways to measure representation quality, freshness, contextual completeness, traceability, and contestability. Second, researchers and practitioners need governance architectures that connect SENSE, CORE, and DRIVER across regulated sectors. Third, enterprises need autonomy allocation methods for deciding when to use deterministic automation, AI agents, human judgment, or hybrid workflows.

Additional research is needed on representation rights. What rights should affected entities have over the representations used to evaluate or act upon them? Market-structure questions also require investigation. What forms of representation monopoly, switching cost, and lock-in may emerge? Security research must address

representation attacks and context manipulation. Finally, standards bodies and industry groups may need common vocabularies, audit practices, and assurance mechanisms for machine-legible institutional reality.

15. Limitations

This article proposes a conceptual framework and does not present empirical validation, controlled experiments, or industry benchmarks. Its purpose is to offer a structured lens for understanding and designing AI-mediated institutional systems. Future research can strengthen the framework by testing it across sectors, developing metrics for representation quality, and evaluating whether SENSE-CORE-DRIVER improves AI governance outcomes in real deployments.

The framework is also not intended to replace existing AI governance standards or laws. It should be read as a practical architecture that can complement them. Its usefulness will depend on disciplined application, sector-specific adaptation, and careful treatment of contested representations.

16. Conclusions

The AI economy is often described as a race for larger models, better agents, faster inference, and cheaper computation. Those races will continue. But they do not fully explain how AI will change enterprises and institutions. AI systems create value only when they can act on a representation of reality that is reliable enough to support judgment and legitimate enough to support action.

The Representation Economy names this shift. It describes an economy in which power accrues to those who can define, maintain, govern, and operationalize machine-legible representations. SENSE-CORE-DRIVER provides an architecture for understanding that shift. SENSE makes reality legible. CORE reasons over that legibility. DRIVER turns decisions into legitimate action.

The core lesson is that better intelligence is not enough. The future of AI will be shaped by the quality of representation and the legitimacy of execution. Organizations that understand this will build AI systems that are not merely capable, but trustworthy, governable, and institutionally useful.

Author Contributions

Conceptualization, R.S.; methodology, R.S.; formal analysis, R.S.; writing - original draft preparation, R.S.; writing - review and editing, R.S.; visualization, R.S. The author has read and agreed to the submitted version of the manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable. This study did not involve humans or animals.

Informed Consent Statement

Not applicable.

Data Availability Statement

No new datasets were generated or analyzed for this conceptual study. Public supporting materials related to the framework are available through the author website and public repository listed in the references.

Acknowledgments

The author acknowledges the broader public discourse on enterprise AI, AI governance, institutional intelligence, and trustworthy AI systems that informed the development of this conceptual framework.

Conflicts of Interest

The author declares no conflicts of interest.

Use of AI-Assisted Tools

AI-assisted tools were used for editorial refinement, formatting support, and language improvement during manuscript preparation. The conceptual framework, terminology, interpretation, final content decisions, and responsibility for the manuscript remain solely with the author.

References

1. Agrawal, A.; Gans, J.; Goldfarb, A. Prediction Machines: The Simple Economics of Artificial Intelligence; Harvard Business Review Press: Boston, MA, USA, 2018.
2. Simon, H.A. The Sciences of the Artificial, 3rd ed.; MIT Press: Cambridge, MA, USA, 1996.
3. Floridi, L. The Fourth Revolution: How the Infosphere Is Reshaping Human Reality; Oxford University Press: Oxford, UK, 2014.
4. Coase, R.H. The Nature of the Firm. Economica 1937, 4, 386-405.
5. Williamson, O.E. The Economic Institutions of Capitalism; Free Press: New York, NY, USA, 1985.
6. Lessig, L. Code and Other Laws of Cyberspace; Basic Books: New York, NY, USA, 1999.
7. Choudary, S.P. Platform Scale; Platform Thinking Labs: Singapore, 2015.
8. Ostrom, E. Governing the Commons: The Evolution of Institutions for Collective Action; Cambridge University Press: Cambridge, UK, 1990.
9. Davenport, T.H.; Ronanki, R. Artificial Intelligence for the Real World. Harvard Business Review 2018, 96, 108-116.
10. Doshi-Velez, F.; Kim, B. Towards a Rigorous Science of Interpretable Machine Learning. arXiv 2017, arXiv:1702.08608.
11. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0); NIST: Gaithersburg, MD, USA, 2023.
12. International Organization for Standardization. ISO/IEC 42001:2023, Information Technology - Artificial Intelligence - Management System; ISO: Geneva, Switzerland, 2023.
13. European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, 2024.
14. OECD. OECD Artificial Intelligence Principles. Available online: https://oecd.ai/en/ai-principles (accessed on 22 May 2026).

15. Porter, T.M. Trust in Numbers: The Pursuit of Objectivity in Science and Public Life; Princeton University Press: Princeton, NJ, USA, 1995.

16. Zuboff, S. The Age of Surveillance Capitalism; PublicAffairs: New York, NY, USA, 2019.

17. Singh, R. SENSE-CORE-DRIVER: The Institutional Architecture That Will Govern the AI Economy. Available online: https://www.raktimsingh.com/sense-core-driver-the-institutional-architecture-that-will-govern-the-ai-economy/ (accessed on 22 May 2026).

18. Singh, R. The Representation Economy: A Framework for AI, Institutions, and Machine-Legible Reality. Zenodo 2026. https://doi.org/10.5281/zenodo.20315480.

19. Singh, R. Representation Economy and SENSE-CORE-DRIVER Public Repository. Available online: https://github.com/raktims2210-dev/representation-economy (accessed on 22 May 2026).

Disclaimer/Publisher Note: The statements and opinions contained in this manuscript are solely those of the author.